Migrating to a new Ubuntu version

Last updated 2024-06-26

To migrate your worker nodes to a new Ubuntu version, you must provision a new worker pool, add worker nodes to the new pool, then remove the original worker pool. The default operating system for all supported IBM Cloud Kubernetes Service cluster versions is Ubuntu 20. Additionally, Ubuntu 24 is available for early access with limitations.

Ubuntu 24 limitations

Important: Ubuntu 24 is available for early release and should not be used for production workloads. Make sure you understand the limitations for this version before you begin any migrations.

Beta: Ubuntu 24 is available in Beta. The following limitations and disclaimers apply.

- Should not be used for production workloads.
- Available for cluster versions 1.29 and later.
- Supported for virtual servers only. Cannot be used with bare metal servers.
- Not available for GPU worker node flavors.
- NTP uses `timesyncd`. Related commands might be updated.
- The following add-ons and features are not supported. Do not migrate your worker nodes if you use these features:
  - CSUtil add-on
  - Object storage plug-in
  - Portworx

Before you begin

1. Review your worker pool operating systems to determine which pools you need to migrate.

```
$ ibmcloud ks worker-pools -c CLUSTER
```

2. For the worker pools that you want to migrate, review the details of the worker pool.

```
$ ibmcloud ks worker-pool get --cluster CLUSTER --worker-pool WORKER-POOL --output json
```

3. In the output, note the zone and either the private and public VLAN ID for Classic clusters or the subnet ID for VPC clusters. Also note any custom labels that you are using. The default worker pool label is `ibm-cloud.kubernetes.io`. Any labels other than `ibm-cloud.kubernetes.io` are custom labels that you should add to your new worker pool.

Migration steps

Migrate your worker nodes to use Ubuntu 24. These steps apply to all supported cluster versions.

1. In your cluster, create a new worker pool for the Ubuntu 24 worker nodes. Include the `--operating-system=UBUNTU_24_64` option. Make sure that the number of nodes specified with the `--size-per-zone` option matches the number of Ubuntu 20 worker nodes that you are replacing. Also, be sure to include the custom labels that you retrieved earlier.


```
$ ibmcloud ks worker-pool create classic --name NAME --cluster CLUSTER --flavor FLAVOR --operating-system
UBUNTU_24_64 --size-per-zone WORKERS-PER-ZONE --label LABEL --label LABEL
```

For VPC clusters. See the CLI reference for command details.

```
$ ibmcloud ks worker-pool create vpc-gen2 --name NAME --cluster CLUSTER --flavor FLAVOR --operating-
system UBUNTU_24_64 --size-per-zone WORKERS-PER-ZONE --label LABEL --label LABEL
```

2. Verify that the worker pool is created.

```
$ ibmcloud ks worker-pool ls --cluster CLUSTER
```

3. Prepare to add a zone to your worker pool. When you add a zone, the number of worker nodes you specified with the `--size-per-zone` option are added to the zone. These worker nodes run the Ubuntu 24 operating system.

Important: Classic and VPC clusters: If possible, use the same Classic VLAN or VPC subnet for the new zone as you are using for the Ubuntu 20 worker pool.

If you need to use a different Classic VLAN or VPC Subnet for these new nodes, note that switching your worker nodes to a different Classic VLAN or VPC subnet can have significant effects your workload and cluster functionality. For example, you might experience the following. - In Classic clusters: Classic LoadBalancers might not work because the LoadBalancer IPs are specific to a single VLAN, and traffic can only be routed to the LoadBalancer if there is a worker node in the cluster on that VLAN that the LoadBalancer IP can be placed on. - In Classic or VPC Clusters: Problems with the network connections to and from workers in the new Classic VLANs or VPC subnets if you have any Security Groups, Network ACLs, Firewall/Gateway rules, Custom routing, coreDNS configurations, and so on, that are specific to the old VLANs, VPC subnets, or cluster worker IP addresses or subnets.

4. Add the zone to your worker pool that you retrieved earlier. When you add a zone, the worker nodes that are defined in your worker pool are provisioned in the zone and considered for future workload scheduling.

- Classic clusters:

```
$ ibmcloud ks zone add classic --zone ZONE --cluster CLUSTER --worker-pool WORKER-POOL --private-vlan
PRIVATE-VLAN-ID --public-vlan PUBLIC-VLAN-ID
```

- VPC clusters:

```
$ ibmcloud ks zone add vpc-gen2 --zone ZONE --cluster CLUSTER --worker-pool WORKER-POOL --subnet-id VPC-
SUBNET-ID
```

5. Verify that worker nodes are available in your new worker pool. In the output, find the listing for the new worker pool and check the number in the **Workers** column.

```
$ ibmcloud ks worker-pool ls --cluster <cluster_name_or-ID>
```

6. **Classic clusters**: If you created your worker pool in a new VLAN, move your ALBs to the new VLAN. For more information, see Moving ALBs across VLANs in classic clusters.

Removing your old worker pools

Important: Before you remove your Ubuntu 20 worker pools, consider scaling them down and keeping them for several days before you remove them. This way, you can scale the worker pool back up if your workload experiences disruptions during the migration process. When you determine that your workload is stable and functions normally, you can remove the Ubuntu 20 worker pool.

1. List your worker pools and note the name of the worker pool you want to remove.

```
$ ibmcloud ks worker-pool ls --cluster CLUSTER
```

2. Run the command to remove the worker pool.

```
$ ibmcloud ks worker-pool rm --worker-pool WORKER-POOL --cluster CLUSTER
```

Contribute in GitHub

Open doc issue Edit topic